

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
8 January 2004 (08.01.2004)

PCT

(10) International Publication Number  
WO 2004/004310 A2

- (51) International Patent Classification: H04N
- (21) International Application Number: PCT/US2003/020397
- (22) International Filing Date: 27 June 2003 (27.06.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Date: 28 June 2002 (28.06.2002) US 10/187,395
- (71) Applicant: DOLBY LABORATORIES LICENSING CORPORATION [US]; 100 Potrero Avenue, San Francisco, CA 94103-4813 (US).
- (72) Inventor: DEMOS, Gary, A.; 10720 Hepburn Circle, Culver City, CA 90232 (US).
- (74) Agent: HARRIS, Scott, C.; Fish & Richardson P.C., Suite 500, 4350 La Jolla Village Drive, San Diego, CA 92122 (US).
- Published: without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guide to the Preparation of International Applications" published by the International Patent Cooperation Treaty.

(54) Title: IMPROVED INTERPOLATION OF VIDEO COMPRESSION FRAMES

(57) Abstract: A method, system, and computer program for improving the image quality of one or more predicted frames in a video image compression system, where each frame comprises a plurality of pixels. A picture region of macroblock of certain types of frames can be encoded by reference to one or more referenceable frames in some cases, and by reference to two or more referenceable frames in other cases. Such encoding may include interpolation, such as an unequal weighting. The DC value or AC pixel values of a picture region may be interpolated as well, with or without weighting. A code pattern of such frames having a variable number of bidirectional predicted frames can be dynamically determined. Frames can be transmitted from an encoder to a decoder in a delivery order different from a display order. Sharpening and/or softening filters can be applied to a picture region of certain frames during motion vector compensated prediction.

# IMPROVED INTERPOLATION OF VIDEO COMPRESSION FRAMES

## CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation-in-part application of U.S. Application serial No. 09/904,203, filed on July 11, 2001 and claims priority to U.S. C.I.P. Application Serial No. 10/187,395 filed June 28, 2002.

## TECHNICAL FIELD

[0002] This invention relates to video compression, and more particularly to improved interpolation of video compression frames in MPEG-like encoding and decoding systems.

## BACKGROUND

### MPEG Video Compression

[0003] MPEG-2 and MPEG-4 are international video compression standards defining respective video syntaxes that provides an efficient way to represent image sequences in the form of more compact coded data. The language of the coded bits is the "syntax." For example, a few tokens can represent an entire block of samples (e.g., 64 samples for MPEG-2). Both MPEG standards also describe a decoding (reconstruction) process where the coded bits are mapped from the compact representation into an approximation of the original format of the image sequence. For example, a flag in the coded bitstream may signal whether the following bits are to be preceded with a prediction algorithm prior to being decoded with a discrete cosine transform (DCT) algorithm. The algorithms comprising the decoding process are regulated by the semantics defined by these MPEG standards. This syntax can be applied to exploit common video characteristics such as spatial redundancy, temporal redundancy, uniform motion, spatial masking, etc. An MPEG decoder must be able to parse and decode an incoming data stream, but so long as the data stream complies with the corresponding MPEG syntax, a wide variety of possible data structures and compression techniques can be used (although technically this deviates from the standard since the semantics are not conformant).

[0004] These MPEG standards use a variety of compression methods, including intraframe and interframe methods. In most video scenes, the background remains relatively stable while action takes place in the foreground. The background may move, but a great deal of the scene often is redundant. These MPEG standards start compression by creating a reference frame called an "intra" frame or "I frame". I frames are compressed without

reference to other frames and thus contain an entire frame of video information. I frames provide entry points into a data bitstream for random access, but can only be moderately compressed. Typically, the data representing I frames is placed in the bitstream every 12 to 15 frames (although it is also useful in some circumstances to use much wider spacing between I frames). Thereafter, since only a small portion of the frames that fall between the reference I frames are different from the bracketing I frames, only the image differences are captured, compressed, and stored. Two types of frames are used for such differences – predicted frames (P frames), and bi-directional predicted (or interpolated) frames (B frames).

[0005] P frames generally are encoded with reference to a past frame (either an I frame or a previous P frame), and, in general, are used as a reference for subsequent P frames. P frames receive a fairly high amount of compression. B frames provide the highest amount of compression but require both a past and a future reference frame in order to be encoded. P and I frames are “referenceable frames” because they can be referenced by P or B frames.

[0006] Macroblocks are regions of image pixels. For MPEG-2, a macroblock is a 16x16 pixel grouping of four 8x8 DCT blocks, together with one motion vector for P frames, and one or two motion vectors for B frames. Macroblocks within P frames may be individually encoded using either intra-frame or inter-frame (predicted) coding. Macroblocks within B frames may be individually encoded using intra-frame coding, forward predicted coding, backward predicted coding, or both forward and backward (i.e., bi-directionally interpolated) predicted coding. A slightly different but similar structure is used in MPEG-4 video coding.

[0007] After coding, an MPEG data bitstream comprises a sequence of I, P, and B frames. A sequence may consist of almost any pattern of I, P, and B frames (there are a few minor semantic restrictions on their placement). However, it is common in industrial practice to have a fixed frame pattern (e.g., IBBPBBPBBPBBPBB).

#### *Motion Vector Prediction*

[0008] In MPEG-2 and MPEG-4 (and similar standards, such as H.263), use of B-type (bi-directionally predicted) frames have proven to benefit compression efficiency. Motion vectors for each macroblock of such frames can be predicted by any one of the following three methods:

[0009] Mode 1: Predicted forward from the previous I or P frame (i.e., a non-bidirectionally predicted frame).

[0010] Mode 2: Predicted backward from the subsequent I or P frame.

[0011] Mode 3: Bi-directionally predicted from both the subsequent and previous I or P frame.

[0012] Mode 1 is identical to the forward prediction method used for P frames. Mode 2 is the same concept, except working backward from a subsequent frame. Mode 3 is an *interpolative* mode that combines information from both previous and subsequent frames.

[0013] In addition to these three modes, MPEG-4 also supports a second interpolative motion vector prediction mode for B frames: *direct mode* prediction using the motion vector from the subsequent P frame, plus a delta value (if the motion vector from the co-located P macroblock is split into 8x8 mode – resulting in four motion vectors for the 16x16 macroblock – then the delta is applied to all four independent motion vectors in the B frame). The subsequent P frame’s motion vector points at the previous P or I frame. A proportion is used to weight the motion vector from the subsequent P frame. The proportion is the relative time position of the current B frame with respect to the subsequent P and previous P (or I) frames.

[0014] FIG. 1 is a time line of frames and MPEG-4 direct mode motion vectors in accordance with the prior art. The concept of MPEG-4 direct mode (mode 4) is that the motion of a macroblock in each intervening B frame is likely to be near the motion that was used to code the same location in the following P frame. A delta is used to make minor corrections to a proportional motion vector derived from the corresponding motion vector (MV) 103 for the subsequent P frame. Shown in FIG. 1 is the proportional weighting given to the motion vectors 101, 102 for each intermediate B frame 104a, 104b as a function of “time distance” between the previous P or I frame 105 and the next P frame 106. The motion vector 101, 102 assigned to a corresponding intermediate B frame 104a, 104b is equal to the assigned weighting value ( $1/3$  and  $2/3$ , respectively) times the motion vector 103 for the next P frame, plus the delta value.

[0015] With MPEG-2, all prediction modes for B frames are tested in coding, and are compared to find the best prediction for each macroblock. If no prediction is good, then the macroblock is coded stand-alone as an “I” (for “intra”) macroblock. The coding mode is selected as the best mode among forward (mode 1), backward (mode 2), and bi-directional (mode 3), or as intra coding. With MPEG-4, the intra coding choice is not allowed. Instead, direct mode becomes the fourth choice. Again, the best coding mode is chosen, based upon some best-match criteria. In the reference MPEG-2 and MPEG-4 software encoders, the best match is determined using a DC match (Sum of Absolute Difference, or “SAD”).

FIG. 1 illustrates the number of successive B frames in a coded data bitstream is determined by the "M" parameter value in MPEG. M minus one is the number of B frames between each P frame and the next P (or I). Thus, for M=3, there are two B frames between each P (or I) frame, as illustrated in FIG. 1. The main limitation in restricting the value of M, and therefore the number of sequential B frames, is that the amount of motion change between P (or I) frames becomes large. Higher numbers of B frames mean longer amounts of time between P (or I) frames. Thus, the efficiency and coding range limitations of motion vectors create the ultimate limit on the number of intermediate B frames.

FIG. 2 illustrates it is also significant to note that P frames carry "change energy" forward with the moving picture stream, since each decoded P frame is used as the starting point to predict the next subsequent P frame. B frames, however, are discarded after use. Thus, any bits used to create B frames are used only for that frame, and do not provide corrections that aid decoding of subsequent frames, unlike P frames.

## SUMMARY

[0018] The invention is directed to a method, system, and computer programs for improving the image quality of one or more predicted frames in a video image compression system, where each frame comprises a plurality of pixels.

[0019] In one aspect, the invention includes determining the value of each pixel of bi-directionally predicted frames as a weighted proportion of corresponding pixel values in non-bidirectionally predicted frames bracketing a sequence of bi-directionally predicted frames.

In one embodiment, the weighted proportion is a function of the distance between the bracketing non-bidirectionally predicted frames. In another embodiment, the weighted proportion is a blended function of the distance between the bracketing non-bidirectionally predicted frames and an equal average of the bracketing non-bidirectionally predicted frames.

[0020] In another aspect of the invention, interpolation of pixel values is performed on representations in a linear space, or in other optimized non-linear spaces differing from an original non-linear representation.

[0021] Other aspects of the invention include systems, computer programs, and methods encompassing:

[0022] • A video image compression system having a sequence of referenceable frames comprising picture regions, in which at least one picture region of at least one predicted frame is encoded by reference to two or more referenceable frames.

[0023] • A video image compression system having a sequence of referenceable frames comprising picture regions, in which at least one picture region of at least one predicted frame is encoded by reference to one or more referenceable frames in display order, where at least one such referenceable frame is not the previous referenceable frame nearest in display order to the at least one predicted frame.

[0024] • A video image compression system having a sequence of referenceable frames comprising macroblocks, in which at least one macroblock within at least one predicted frame is encoded by interpolation from two or more referenceable frames.

[0025] • A video image compression system having a sequence of referenceable and bidirectional predicted frames comprising picture regions, in which at least one picture region of at least one bidirectional predicted

frame is encoded to include more than two motion vectors, each such motion vector referencing a corresponding picture region in at least one referenceable frame.

- [0026] • A video image compression system having a sequence of referenceable frames comprising picture regions, in which at least one picture region of at least one predicted frame is encoded to include at least two motion vectors, each such motion vector referencing a corresponding picture region in a referenceable frame, where each such picture region of such at least one predicted frame is encoded by interpolation from two or more referenceable frames.
- [0027] • A video image compression system having a sequence of referenceable and bidirectional predicted frames comprising picture regions, in which at least one picture region of at least one bidirectional predicted frame is encoded as an unequal weighting of selected picture regions from two or more referenceable frames.
- [0028] • A video image compression system having a sequence of referenceable and bidirectional predicted frames comprising picture regions, in which at least one picture region of at least one bidirectional predicted frame is encoded by interpolation from two or more referenceable frames, where at least one of the two or more referenceable frames is spaced from the bidirectional predicted frame by at least one intervening referenceable frame in display order, and where such at least one picture region is encoded as an unequal weighting of selected picture regions of such at least two or more referenceable frames.
- [0029] • A video image compression system having a sequence of referenceable and bidirectional predicted frames comprising picture regions, in which at least one picture region of at least one bidirectional predicted frame is encoded by interpolation from two or more referenceable frames, where at least one of the two or more referenceable frames is spaced from the bidirectional predicted frame by at least one intervening subsequent referenceable frame in display order.
- [0030] • A video image compression system having a sequence of referenceable and bidirectional predicted frames comprising picture regions, in

which at least one picture region of at least one bidirectional predicted frame is encoded as an unequal weighting from selected picture regions of two or more referenceable frames.

- [0031] • A video image compression system having a sequence of predicted and bidirectional predicted frames each comprising pixel values arranged in macroblocks, wherein at least one macroblock within a bidirectional predicted frame is determined using direct mode prediction based on motion vectors from two or more predicted frames.
- [0032] • A video image compression system having a sequence of referenceable and bidirectional predicted frames each comprising pixel values arranged in macroblocks, wherein at least one macroblock within a bidirectional predicted frame is determined using direct mode prediction based on motion vectors from one or more predicted frames in display order, wherein at least one of such one or more predicted frames is previous in display order to the bidirectional predicted frame.
- [0033] • A video image compression system having a sequence of referenceable and bidirectional predicted frames each comprising pixel values arranged in macroblocks, wherein at least one macroblock within a bidirectional predicted frame is determined using direct mode prediction based on motion vectors from one or more predicted frames, wherein at least one of such one or more predicted frames is subsequent in display order to the bidirectional predicted frame and spaced from the bidirectional predicted frame by at least one intervening referenceable frame.
- [0034] • A video image compression system having a sequence of frames comprising a plurality of picture regions having a DC value, each such picture region comprising pixels each having an AC pixel value, wherein at least one of the DC value and the AC pixel values of at least one picture region of at least one frame are determined as a weighted interpolation of corresponding respective DC values and AC pixel values from at least one other frame.
- [0035] • A video image compression system having a sequence of referenceable frames comprising a plurality of picture regions having a DC value,

each such picture region comprising pixels each having an AC pixel value, in which at least one of the DC value and AC pixel values of at least one picture region of at least one predicted frame are interpolated from corresponding respective DC values and AC pixel values of two or more referenceable frames.

- [0036] Improving the image quality of a sequence of two or more bidirectional predicted intermediate frames in a video image compression system, each frame comprising a plurality picture regions having a DC value, each such picture region comprising pixels each having an AC pixel value, including at least one of the following: determining the AC pixel values of each picture region of a bidirectional predicted intermediate frame as a first weighted proportion of corresponding AC pixel values in referenceable frames bracketing the sequence of bidirectionally predicted intermediate frames; and determining the DC value of each picture region of such bidirectional predicted intermediate frame as a second weighted proportion of corresponding DC values in referenceable frames bracketing the sequence of bidirectional predicted intermediate frames.

A video image compression system having a sequence of frames comprising a plurality of pixels having an initial representation, in which the pixels of at least one frame are interpolated from corresponding pixels of at least two other frames, wherein such corresponding pixels of the at least two other frames are interpolated while transformed to a different representation, and the resulting interpolated pixels are transformed back to the initial representation.

- [0037] In a video image compression system having a sequence of referenceable and bidirectional predicted frames, dynamically determining a code pattern of such frames having a variable number of bidirectional predicted frames, including: selecting an initial sequence beginning with a referenceable frame, having at least one immediately subsequent bidirectional predicted frame, and ending in a referenceable frame; adding a referenceable frame to the end of the initial sequence to create a test sequence; evaluating the test sequence against a

selected evaluation criteria; for each satisfactory step of evaluating the test sequence, inserting a bidirectional frame before the added referenceable frame and repeating the step of evaluating; and if evaluating the test sequence is unsatisfactory, then accepting the prior test sequence as a current code pattern.

- [0038] A video image compression system having a sequence of referenceable frames spaced by at least one bidirectional predicted frames, wherein the number of such bidirectional predicted frames varies in such sequence, and wherein at least one picture region of at least one such bidirectional predicted frame is determined using an unequal weighting of pixel values corresponding to at least two referenceable frames.
- [0039] A video image compression system having a sequence of frames encoded by a coder for decoding by a decoder, wherein at least one picture region of at least one frame is based on weighted interpolations of two or more other frames, such weighted interpolations being based on at least one set of weights available to the coder and a decoder, wherein a designation for a selected one of such at least one set of weights is communicated to a decoder from the coder to select one or more currently active weights.

- [0040] A video image compression system having a sequence of frames encoded by a coder for decoding by a decoder, wherein at least one picture region of at least one frame is based on weighted interpolations of two or more other frames, such weighted interpolations being based on at least one set of weights, wherein at least one set of weights is downloaded to a decoder and thereafter a designation for a selected one of such at least one set of weights is communicated to a decoder from the coder to select one or more currently active weights.

- [0041] A video image compression system having a sequence of referenceable frames encoded by a coder for decoding by a decoder, wherein predicted frames in the sequence of referenceable frames are transmitted by the encoder to the decoder in a delivery order that differs from the display order of such predicted frames after decoding.

- [0042] • A video image compression system having a sequence of referenceable frames comprising pixels arranged in picture regions, in which at least one picture region of at least one predicted frame is encoded by reference to two or more referenceable frames, wherein each such picture region is determined using an unequal weighting of pixel values corresponding to such two or more referenceable frames.
- [0043] • A video image compression system having a sequence of predicted, bidirectional predicted, and intra frames each comprising picture regions, wherein at least one filter selected from the set of sharpening and softening filters is applied to at least one picture region of a predicted or bidirectional predicted frame during motion vector compensated prediction of such predicted or bidirectional predicted frame.
- [0044] The details of one or more embodiments of the invention are set forth in the accompanying drawings and the description below. Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

#### DESCRIPTION OF DRAWINGS

- [0045] FIG. 1 is a time line of frames and MPEG-4 direct mode motion vectors in accordance with the prior art.
- [0046] FIG. 2 is a time line of frames and proportional pixel weighting values in accordance with this aspect of the invention.
- [0047] FIG. 3 is a time line of frames and blended, proportional, and equal pixel weighting values in accordance with this aspect of the invention.
- [0048] FIG. 4 is a flowchart showing an illustrative embodiment of the invention as a method that may be computer implemented.
- [0049] FIG. 5 is a diagram showing an example of multiple previous references by a current P frame to two prior P frames, and to a prior I frame.
- [0050] FIG. 6A is a diagram of a typical prior art MPEG-2 coding pattern, showing a constant number of B frames between bracketing I frames and/or P frames.
- [0051] FIG. 6B is a diagram of a theoretically possible prior art MPEG-4 video coding pattern, showing a varying number of B frames between bracketing I frames and/or P frames, as well as a varying distance between I frames.
- [0052] FIG. 7 is a diagram of code patterns.

- [0053] FIG. 8 is a flowchart showing one embodiment of an interpolation method with DC interpolation being distinct from AC interpolation.
- [0054] FIG. 9 is a flowchart showing one embodiment of a method for interpolation of luminance pixels using an alternative representation.
- [0055] FIG. 10 is a flowchart showing one embodiment of a method for interpolation of chroma pixels using an alternative representation.
- [0056] FIG. 11 is a diagram showing unique motion vector region sizes for each of two P frames.
- [0057] FIG. 12 is a diagram showing a sequence of P and B frames with interpolation weights for the B frames determined as a function of distance from a 2-away subsequent P frame that references a 1-away subsequent P frame.
- [0058] FIG. 13 is a diagram showing a sequence of P and B frames with interpolation weights for the B frames determined as a function of distance from a 1-away subsequent P frame that references a 2-away previous P frame.
- [0059] FIG. 14 is a diagram showing a sequence of P and B frames in which a subsequent P frame has multiple motion vectors referencing prior P frames.
- [0060] FIG. 15 is a diagram showing a sequence of P and B frames in which a nearest subsequent P frame has a motion vector referencing a prior P frame, and a next nearest subsequent P frame has multiple motion vectors referencing prior P frames.
- [0061] FIG. 16 is a diagram showing a sequence of P and B frames in which a nearest previous P frame has a motion vector referencing a prior P frame.
- [0062] FIG. 17 is a diagram showing a sequence of P and B frames in which a nearest previous P frame has two motion vectors referencing prior P frames.
- [0063] FIG. 18 is a diagram showing a sequence of P and B frames in which a nearest previous P frame has a motion vector referencing a prior P frame.
- [0064] FIG. 19 is a frame sequence showing the case of three P frames P1, P2, and P3, where P3 uses an interpolated reference with two motion vectors, one for each of P1 and P2.
- [0065] FIG. 20 is a frame sequence showing the case of four P frames P1, P2, P3, and P4, where P4 uses an interpolated reference with three motion vectors, one for each of P1, P2, and P3.
- [0066] FIG. 21 is a diagram showing a sequence of P and B frames in which various P frames have one or more motion vectors referencing various previous P frames, and

showing different weights assigned to respective forward and backward references by a particular B frame.

[0067] FIG. 22 is a diagram showing a sequence of P and B frames in which the bitstream order of the P frames differs from the display order.

[0068] FIG. 23 is a diagram showing a sequence of P and B frames with assigned weightings.

[0069] FIG. 24 is a graph of position of an object within a frame versus time.

[0070] Like reference symbols in the various drawings indicate like elements.

## DETAILED DESCRIPTION

### Overview

[0071] One aspect of the invention is based upon recognition that it is common practice to use a value for  $M$  of 3, which provides for two B frames between each P (or I) frame. However  $M=2$ , and  $M=4$  or higher, are all useful. It is of particular significance to note that the value of  $M$  (the number of B frames plus 1) also bears a natural relationship to the frame rate. At 24 frames per second (fps), the rate of film movies, the 1/24th second time distance between frames can result in substantial image changes frame-to-frame. At 60 fps, 72 fps, or higher frame rates, however, the time distance between adjacent frames becomes correspondingly reduced. The result is that higher numbers of B frames (*i.e.*, higher values of  $M$ ) become useful and beneficial in compression efficiency as the frame rate is increased.

[0072] Another aspect of the invention is based upon the recognition that both MPEG-2 and MPEG-4 video compression utilize an oversimplified method of interpolation. For example, for mode 3, the bi-directional prediction for each macroblock of a frame is an equal average of the subsequent and previous frame macroblocks, as displaced by the two corresponding motion vectors. This equal average is appropriate for  $M=2$  (*i.e.*, single intermediate B frames), since the B frame will be equidistant in time from the previous and subsequent P (or I) frames. However, for all higher values of  $M$ , only symmetrically centered B frames (*i.e.*, the middle frame if  $M=4$ , 6, 8, *etc.*) will be optimally predicted using an equal weighting. Similarly, in MPEG-4 direct mode 4, even though the *motion vectors* are proportionally weighted, the predicted *pixel values* for each intermediate B frame are an equal proportion of the corresponding pixels of the previous P (or I) and subsequent P frame.

[0073] Thus, it represents an improvement to apply an appropriate proportional weighting, for  $M>2$ , to the predicted pixel values for each B frame. The proportional weighting for each pixel in a current B frame corresponds to the relative position of the current B frame with respect to the previous and subsequent P (or I) frames. Thus, if  $M=3$ , the first B frame would use 2/3 of the corresponding pixel value (motion vector adjusted) from the previous frame, and 1/3 of the corresponding pixel value from the subsequent frame (motion vector adjusted).

[0074] FIG. 2 is a time line of frames and proportional pixel weighting values in accordance with this aspect of the invention. The pixel values within each macroblock of each intermediate B frame 201a, 201b are weighted as a function of "distance" between the previous P or I frame *A* and the next P or I frame *B*, with greater weight being accorded to

closer I or P frames. That is, each pixel value of a bi-directionally predicted B frame is a weighted combination of the corresponding pixel values of bracketing non-bidirectionally predicted frames A and B. In this example, for  $M=3$ , the weighting for the first B frame 201a is equal to  $2/3A + 1/3B$ ; the weighting for the second B frame 201b is equal to  $1/3A + 2/3B$ . Also shown is the equal average weighting that would be assigned under conventional MPEG systems; the MPEG-1, 2, and 4 weighting for each B frame 201a, 201b would be equal to  $(A + B)/2$ .

*Application to Extended Dynamic Range and Contrast Range*

[0075] If M is greater than 2, proportional weighting of pixel values in intermediate B frames will improve the effectiveness of bi-directional (mode 3) and direct (MPEG-4 mode 4) coding in many cases. Example cases include common movie and video editing effects such as fade-outs and cross-dissolves. These types of video effects are problem coding cases for both MPEG-2 and MPEG-4 due to use of a simple DC matching algorithm, and the common use of  $M=3$  (i.e., two intermediate B frames), resulting in equal proportions for B frames. Coding of such cases is improved by using proportional B frame interpolation in accordance with the invention.

[0076] Proportional B frame interpolation also has direct application to coding efficiency improvement for extending dynamic and contrast range. A common occurrence in image coding is a change in illumination. This occurs when an object moves gradually into (or out from) shadow (soft shadow edges). If a logarithmic coding representation is used for brightness (as embodied by logarithmic luminance Y, for example), then a lighting brightness change will be a DC offset change. If the brightness of the lighting drops to half, the pixel values will all be decreased by an equal amount. Thus, to code this change, an AC match should be found, and a coded DC difference applied to the region. Such a DC difference being coded into a P frame should be proportionally applied in each intervening B frame as well. (See co-pending U.S. Patent Application No. 09/905,039, entitled "Method and System for Improving Compressed Image Chroma Information", assigned to the assignee of the present invention and hereby incorporated by reference, for additional information on logarithmic coding representations).

[0077] In addition to changes in illumination, changes in contrast also benefit from proportional B frame interpolation. For example, as an airplane moves toward a viewer out of a cloud or haze, its contrast will gradually increase. This contrast increase will be expressed as an increased amplitude in the AC coefficients of the DCT in the corresponding P frame

coded macroblocks. Again, contrast changes in intervening B frames will be most closely approximated by a proportional interpolation, thus improving coding efficiency.

[0078] Improvements in dynamic range and contrast coding efficiency using proportional B frame interpolation become increasingly significant as frame rates become higher and as the value of M is increased.

*Applying High M Values to Temporal Layering*

[0079] Using embodiments of the invention allows an increase in the value of M, and hence the number of B frames between bracketing P and/or I frames, while maintaining or gaining coding efficiency. Such usage benefits a number of applications, including temporal layering. For example, in U.S. Patent No. 5,988,863, entitled "Temporal and Resolution Layering for Advanced Television" (assigned to the assignee of the present invention, and incorporated by reference), it was noted that B frames are a suitable mechanism for layered temporal (frame) rates. The flexibility of such rates is related to the number of consecutive B frames available. For example, single B frames ( $M=2$ ) can support a 36 fps decoded temporal layer within a 72 fps stream or a 30 fps decoded temporal layer within a 60 fps stream. Triple B frames ( $M=4$ ) can support both 36 fps and 18 fps decoded temporal layers within a 72 fps stream, and 30 fps and 15 fps decoded temporal layers within a 60 fps stream. Using  $M=10$  within a 120 fps stream can support 12 fps, 24 fps, and 60 fps decoded temporal layers.  $M=4$  also can be used with a 144 fps stream to provide for decoded temporal layers at 72 fps and 36 fps.

[0080] As an improvement to taking every  $N^{\text{th}}$  frame, multiple frames at 120 fps or 72 fps can be decoded and proportionally blended, as described in co-pending U.S. Patent Application No. 09/545,233, entitled "Enhancements to Temporal and Resolution Layering" (assigned to the assignee of the present invention and incorporated by reference), to improve the motion blur characteristics of the 24 fps results.

[0081] Even higher frame rates can be synthesized utilizing the methods described in co-pending U.S. Patent Application No. 09/435,277, entitled "System and Method for Motion Compensation and Frame Rate Conversion" (assigned to the assignee of the present invention and incorporated by reference). For example, a 72 fps camera original can be utilized with motion compensated frame rate conversion to create an effective frame rate of 288 frames per second. Using  $M=12$ , both 48 fps and 24 fps frame rates can be derived, as well as other useful rates such as 144 fps, 96 fps, and 32 fps (and of course, the original 72 fps). The frame rate conversions using this method need not be integral multiples. For example, an effective



rate of 120 fps can be created from a 72 fps source, and then used as a source for both 60 fps and 24 fps rates (using  $M=10$ ).

[0082] Thus, there are temporal layering benefits to optimizing the performance of B frame interpolation. The proportional B frame interpolation described above make higher numbers of consecutive B frames function more efficiently, thereby enabling these benefits.

#### *Blended B-Frame Interpolation Proportions*

[0083] One reason that equal average weighting has been used in conventional systems as the motion compensated mode predictor for B frame pixel values is that the P (or I) frame before or after a particular B frame may be noisy, and therefore represent an imperfect match. Equal blending will optimize the reduction of noise in the interpolated motion-compensated block. There is a difference residual that is coded using the quantized DCT function. Of course, the better the match from the motion compensated proportion, the fewer difference residual bits will be required, and the higher the resulting image quality.

[0084] In cases where there are objects moving in and out of shadow or haze, a true proportion where  $M>2$  provides a better prediction. However, when lighting and contrast changes are not occurring, equal weighting may prove to be a better predictor, since the errors of moving a macroblock forward along a motion vector will be averaged with the errors from the backward displaced block, thus reducing the errors in each by half. Even so, it is more likely that B frame macroblocks nearer a P (or I) frame will correlate more to that frame than to a more distant P (or I) frame.

[0085] Thus, it is desirable in some circumstances, such as regional contrast or brightness change, to utilize a true proportion for B frame macroblock pixel weighting (for both luminance and color), as described above. In other circumstances, it may be more optimal to utilize equal proportions, as in MPEG-2 and MPEG-4.

[0086] Another aspect of the invention utilizes a blend of these two proportion techniques (equal average and frame-distance proportion) for B frame pixel interpolation. For example, in the  $M=3$  case,  $3/4$  of the  $1/3$  and  $2/3$  proportions can be blended with  $1/4$  of the equal average, resulting in the two proportions being  $3/8$  and  $5/8$ . This technique may be generalized by using a "blend factor" F:

$$Weight = F \cdot (FrameDistanceProportionalWeight) + (1-F) \cdot (EqualAverageWeight)$$

The useful range of the blend factor F is from 1, indicating purely proportional interpolation, to 0, indicating purely equal average (the reverse assignment of values may also be used).

[0087] FIG. 3 is a time line of frames and blended, proportional, and equal pixel weighting values in accordance with this aspect of the invention. The pixel values of each macroblock of each intermediate B frame 301a, 301b are weighted as a function of "time distance" between the previous P or I frame A and the next P or I frame B, and as a function of the equal average of A and B. In this example, for  $M=3$  and a blend factor  $F=3/4$ , the blended weighting for the first B frame 301a is equal to  $5/8A + 3/8B$  (i.e.,  $3/4$  of the proportional weighting of  $2/3A + 1/3B$ , plus  $1/4$  of the equal average weighting of  $(A + B)/2$ ). Similarly, the weighting for the second B frame 301b is equal to  $3/8A + 5/8B$ .

[0088] The value of the blend factor F can be set overall for a complete encoding, or for each group of pictures (GOP), a range of B frames, each B frame, or each region within a B frame (including, for example, as finely as for each macroblock or, in the case of MPEG-4 direct mode using a P vector in 8x8 mode, even individual 8x8 motion blocks).

[0089] In the interest of bit economy, and reflecting the fact that the blend proportion is not usually important enough to be conveyed with each macroblock, optimal use of blending should be related to the type of images being compressed. For example, for images that are fading, dissolving, or where overall lighting or contrast is gradually changing, a blend factor F near or at 1 (i.e., selecting proportional interpolation) is generally most optimal. For running images without such lighting or contrast changes, then lower blend factor values, such as  $2/3$ ,  $1/2$ , or  $1/3$ , might form a best choice, thereby preserving some of the benefits of proportional interpolation as well as some of the benefits of equal average interpolation. All blend factor values within the 0 to 1 range generally will be useful, with one particular value within this range proving optimal for any given B frame.

[0090] For wide dynamic range and wide contrast range images, the blend factor can be determined regionally, depending upon the local region characteristics. In general, however, a wide range of light and contrast recommends toward blend factor values favoring purely proportional, rather than equal average, interpolation.

[0091] An optimal blend factor is generally empirically determined, although experience with particular types of scenes can be used to create a table of blend factors by scene type. For example, a determination of image change characteristics can be used to select the blend proportion for a frame or region. Alternatively, B frames can be coded using a number of candidate blend factors (either for the whole frame, or regionally), with each then being evaluated to optimize the image quality (determined, for example, by the highest signal to noise ratio, or SNR) and for lowest bit count. These candidate evaluations can then

be used to select the best value for the blend proportion. A combination of both image change characteristics and coded quality/efficiency can also be used.

[0092] B frames near the middle of a sequence of B frames, or resulting from low values of M, are not affected very much by proportional interpolation, since the computed proportions are already near the equal average. However, for higher values of M, the extreme B frame positions can be significantly affected by the choice of blend factor. Note that the blend factor can be different for these extreme positions, utilizing more of the average, than the more central positions, which gain little or no benefit from deviating from the average, since they already have high proportions of both neighboring P (or I) frames. For example, if  $M=5$ , the first and fourth B frame might use a blend factor F which blends in more of the equal average, but the second and third middle B frames may use the strict  $2/5$  and  $3/5$  equal average proportions. If the proportion-to-average blend factor varies as a function of the position of a B frame in a sequence, the varying value of the blend factor can be conveyed in the compressed bitstream or as side information to the decoder.

[0093] If a static general blend factor is required (due to lack of a method to convey the value), then the value of  $2/3$  is usually near optimal, and can be selected as a static value for B frame interpolation in both the encoder and decoder. For example, using  $F=2/3$  for the blend factor, for  $M=3$  the successive frame proportions will be  $7/18 = 2/3 * 1/3 + 1/3 * 1/2$  and  $11/18 = 2/3 * 2/3 + 1/3 * 1/2$ .

#### Linear Interpolation

[0094] Video frame pixel values are generally stored in a particular representation that maps the original image information to numeric values. Such a mapping may result in a linear or non-linear representation. For example, luminance values used in compression are non-linear. The use of various forms of non-linear representation include logarithmic, exponential (to various powers), and exponential with a black correction (commonly used for video signals).

[0095] Over narrow dynamic ranges, or for interpolations of nearby regions, the non-linear representation is acceptable, since these nearby interpolations represent piece-wise linear interpolations. Thus, small variations in brightness are reasonably approximated by linear interpolation. However, for wide variations in brightness, such as occur in wide dynamic range and wide contrast range images, the treatment of non-linear signals as linear will be inaccurate. Even for normal contrast range images, linear fades and cross-dissolves

can be degraded by a linear interpolation. Some fades and cross-dissolves utilize non-linear fade and dissolve rates, adding further complexity.

[0096] Thus, an additional improvement to the use of proportional blends, or even simple proportional or equal average interpolations, is to perform such interpolations on pixel values represented in a linear space, or in other optimized non-linear spaces differing from the original non-linear luminance representation.

[0097] This may be accomplished, for example, by first converting the two non-linear luminance signals (from the previous and subsequent P (or I) frames into a linear representation, or a differing non-linear representation. Then a proportional blend is applied, after which the inverse conversion is applied, yielding the blended result in the image's original non-linear luminance representation. However, the proportion function will have been performed on a more optimal representation of the luminance signals.

[0098] It is also useful to beneficially apply this linear or non-linear conversion to color (chroma) values, in addition to luminance, when colors are fading or becoming more saturated, as occurs in contrast changes associated with variations in haze and overcast.

#### Example Embodiment

[0099] FIG. 4 is a flowchart showing an illustrative embodiment of the invention as a method that may be computer implemented:

[00100] Step 400: In a video image compression system, for direct and interpolative mode for computing B frames, determine an interpolation value to apply to each pixel of an input sequence of two or more bi-directionally predicted intermediate frames using one of (1) the frame-distance proportion or (2) a blend of equal weighting and the frame-distance proportion, derived from at least two non-bidirectionally predicted frames bracketing such sequence input from a source (e.g., a video image stream).

[00101] Step 401: Optimize the interpolation value with respect to an image unit (e.g., a group of pictures (GOP), a sequence of frames, a scene, a frame, a region within a frame, a macroblock, a DCT block, or similar useful grouping or selection of pixels). The interpolation value may be set statically for the entire encoding session, or dynamically for each image unit.

[00102] Step 402: Further optimize the interpolation value with respect to scene type or coding simplicity. For example, an interpolation value may be set: statically (such as  $2/3$  proportional and  $1/3$  equal average); proportionally for frames near the equal average, but blended with equal average near the adjacent P (or I) frames; dynamically based upon overall

scene characteristics, such as fades and cross dissolves; dynamically (and locally) based on local image region characteristics, such as local contrast and local dynamic range; or dynamically (and locally) based upon coding performance (such as highest coded SNR) and minimum coded bits generated.

[00103] Step 403: Convey the appropriate proportion amounts to the decoder, if not statically determined.

[00104] Step 404: Optionally, convert the luminance (and, optionally, chroma) information for each frame to a linear or alternate non-linear representation, and convey this alternate representation to the decoder, if not statically determined.

[00105] Step 405: Determine the proportional pixel values using the determined interpolation value.

[00106] Step 406: If necessary (because of Step 404), reconvert to the original representation.

#### *Extended P frame reference*

[00107] As noted above, in prior art MPEG-1, 2, and 4 compression methods, P frames reference the previous P or I frame, and B frames reference the nearest previous and subsequent P and/or I frames. The same technique is used in the H.261 and H.263 motion-compensated DCT compression standards, which encompass low bit rate compression techniques.

[00108] In the H.263++ and H.26L standard in development, B frame referencing was extended to point to P or I frames which were not directly bracketing a current frame. That is, macro blocks within B frames could point to one P or I frame before the previous P frame, or to one P or I frame after the subsequent P frame. With one or more bits per macroblock, skipping of the previous or subsequent P frame can be signaled simply. Conceptually, the use of previous P frames for reference in B frames only requires storage. For the low-bit rate-coding use of H.263++ or H.26L, this is a small amount of additional memory. For subsequent P frame reference, the P frame coding order must be modified with respect to B frame coding, such that future P frames (or possibly I frames) must be decoded before intervening B frames. Thus, coding order is also an issue for subsequent P frame references.

[00109] The primary distinctions between P and B frame types are: (1) B frames may be bi-directionally referenced (up to two motion vectors per macroblock); (2) B frames are discarded after use (which also means that they can be skipped during decoding to provide

temporal layering); and (3) P frames are used as "stepping stones", one to the next, since each P frame must be decoded for use as a reference for each subsequent P frame.

[00110] As another aspect of the invention, P frames (as opposed to B frames) are decoded with reference to one or more previous P or I frames (excluding the case of each P frame referencing only the nearest previous P or I frame). Thus, for example, two or more motion vectors per macroblock may be used for a current P frame, all pointing backward in time (i.e., to one or more previously decoded frames). Such P frames still maintain a "stepping stone" character. FIG. 5 is a diagram showing an example of multiple previous references by a current P frame 500 to two prior P frames 502, 504, and to a prior I frame 506.

[00111] Further, it is possible to apply the concepts of macroblock interpolation, as described above, in such P frame references. Thus, in addition to signaling single references to more than one previous P or I frame, it is also possible to blend proportions of multiple previous P or I frames, using one motion vector for each such frame reference. For example, the technique described above of using a B frame interpolation mode having two frame references may be applied to allow any macroblock in a P frame to reference two previous P frames or one previous P frame and one previous I frame, using two motion vectors. This technique interpolates between two motion vectors, but is not bi-directional in time (as is the case with B frame interpolation), since both motion vectors point backward in time. Memory costs have decreased to the point where holding multiple previous P or I frames in memory for such concurrent reference is quite practical.

[00112] In applying such P frame interpolation, it is constructive to select and signal to a decoder various useful proportions of the previous two or more P frames (and, optionally, one prior I frame). In particular, an equal blend of frames is one of the useful blend proportions. For example, with two previous P frames as references, an equal 1/2 amount of each P frame can be blended. For three previous P frames, a 1/3 equal blend could be used.

[00113] Another useful blend of two P frames is 2/3 of the most recent previous frame, and 1/3 of the least recent previous frame. For three previous P frames, another useful blend is 1/2 of the most recent previous frame, 1/3 of the next most recent previous frame, and 1/6 of the least recent previous frame.

[00114] In any case, a simple set of useful blends of multiple previous P frames (and, optionally, one I frame) can be utilized and signaled simply from an encoder to a decoder. The specific blend proportions utilized can be selected as often as useful to optimize coding efficiency for an image unit. A number of blend proportions can be selected using a small

number of bits, which can be conveyed to the decoder whenever suitable for a desired image unit.

[00115] As another aspect of the invention, it is also useful to switch-select single P frame references from the most recent previous P (or I) frame to a more "distant" previous P (or I) frame. In this way, P frames would utilize a single motion vector per macroblock (or, optionally, per 8x8 block in MPEG-4 style coding), but would utilize one or more bits to indicate that the reference refers to a single specific previous frame. P frame macroblocks in this mode would not be interpolative, but instead would reference a selected previous frame, being selected from a possible two, three, or more previous P (or I) frame choices for reference. For example, a 2-bit code could designate one of up to four previous frames as the single reference frame of choice. This 2-bit code could be changed at any convenient image unit.

#### *Adaptive Number of B Frames*

[00116] It is typical in MPEG coding to use a fixed pattern of I, P, and B frame types. The number of B frames between P frames is typically a constant. For example, it is typical in MPEG-2 coding to use two B frames between P (or I) frames. FIG. 6A is a diagram of a typical prior art MPEG-2 coding pattern, showing a constant number of B frames (i.e., two) between bracketing I frames 600 and/or P frames 602.

[00117] The MPEG-4 video coding standard conceptually allows a varying number of B frames between bracketing I frames and/or P frames, and a varying amount of distance between I frames. FIG. 6B is a diagram of a theoretically possible prior art MPEG-4 video coding pattern, showing a varying number of B frames between bracketing I frames 600 and/or P frames 602, as well as a varying distance between I frames 600.

[00118] This flexible coding structure theoretically can be utilized to improve coding efficiency by matching the most effective B and P frame coding types to the moving image frames. While this flexibility has been specifically allowed, it has been explored very little, and no mechanism is known for actually determining the placement of B and P frames in such a flexible structure.

[00119] Another aspect of the invention applies the concepts described herein to this flexible coding structure as well as to the simple fixed coding patterns in common use. B frames thus can be interpolated using the methods described above, while P frames may reference more than one previous P or I frame and be interpolated in accordance with the present description.

[00120] In particular, macroblocks within B frames can utilize proportional blends appropriate for a flexible coding structure as effectively as with a fixed structure.

Proportional blends can also be utilized when B frames reference P or I frames that are further away than the nearest bracketing P or I frames.

[00121] Similarly, P frames can reference more than one previous P or I frame in this flexible coding structure as effectively as with a fixed pattern structure. Further, blend proportions can be applied to macroblocks in such P frames when they reference more than one previous P frame (plus, optionally, one I frame).

#### *(A) Determining Placement in Flexible Coding Patterns*

[00122] The following method allows an encoder to optimize the efficiency of both the frame coding pattern as well as the blend proportions utilized. For a selected range of frames, a number of candidate coding patterns can be tried, to determine an optimal or near optimal (relative to specified criteria) pattern. FIG. 7 is a diagram of code patterns that can be examined. An initial sequence 700, ending in a P or I frame, is arbitrarily selected, and is used as a base for adding additional P and/or B frames, which are then evaluated (as described below). In one embodiment, a P frame is added to the initial sequence 700 to create a first test sequence 702 for evaluation. If the evaluation is satisfactory, an intervening B frame is inserted to create a second test sequence 704. For each satisfactory evaluation, additional B frames are inserted to create increasingly longer test sequences 706-712, until the evaluation criteria become unsatisfactory. At that point, the previous coding sequence is accepted. This process is then repeated, using the end P frame for the previously accepted coding sequence as the starting point for adding a new P frame and then inserting new B frames.

[00123] An optimal or near optimal coding pattern can be selected based upon various evaluation criteria, often involving tradeoffs of various coding characteristics, such as coded image quality versus number of coding bits required. Common evaluation criteria include the least number of bits used (in a fixed quantization parameter test), or the best signal-to-noise ratio (in a fixed bit-rate test), or a combination of both.

[00124] It is also common to minimize a sum-of-absolute-difference (SAD), which forms a measure of DC match. As described in co-pending U.S. Patent No. 09/904,192, entitled "Motion Estimation for Video Compression Systems" (assigned to the assignee of the present invention and hereby incorporated by reference), an AC match criterion is also a useful measure of the quality of a particular candidate match (the patent application also

describes other useful optimizations). Thus, the AC and DC match criteria, accumulated over the best matches of all macroblocks, can be examined to determine the overall match quality of each candidate coding pattern. This AC/DC match technique can augment or replace the signal-to-noise ratio (SNR) and least-bits-used tests when used together with an estimate of the number of coded bits for each frame pattern type. It is typical to code macroblocks within B frames with a higher quantization parameter (QP) value than for P frames, affecting both the quality (measured often as a signal-to-noise ratio) and the number of bits used within the various candidate coding patterns.

#### (B) Blend Proportion Optimization in Flexible Coding Patterns

[00125] Optionally, for each candidate pattern determined in accordance with the above method, blend proportions may be tested for suitability (e.g., optimal or near optimal blend proportions) relative to one or more criteria. This can be done, for example, by testing for best quality (lowest SNR) and/or efficiency (least bits used). The use of one or more previous references for each macroblock in P frames can also be determined in the same way, testing each candidate reference pattern and blend proportion, to determine a set of one or more suitable references.

[00126] Once the coding pattern for this next step (Step 700 in FIG. 7) has been selected, then the subsequent steps (Steps 702-712) can be tested for various candidate coding patterns. In this way, a more efficient coding of a moving image sequence can be determined. Thus, efficiency can be optimized/improved as described in subsection (A) above; blend optimization can be applied at each tested coding step.

#### DC vs. AC Interpolation

[00127] In many cases of image coding, such as when using a logarithmic representation of image frames, the above-described interpolation of frame pixel values will optimally code changes in illumination. However, in alternative video "gamma-curve", linear, and other representations, it will often prove useful to apply different interpolation blend factors to the DC values than to the AC values of the pixels. FIG. 8 is a flowchart showing one embodiment of an interpolation method with DC interpolation being distinct from AC interpolation. For a selected image region (usually a DCT block or macroblock) from a first and second input frame 802, 802', the average pixel value for each such region is subtracted 804, 804', thereby separating the DC value (i.e., the average value of the entire selected region) 806, 806' from the AC values (i.e., the signed pixel values remaining) 808,

808' in the selected regions. The respective DC values 806, 806' can then be multiplied by interpolation weightings 810, 810' different from the interpolation weightings 814, 814' used to multiply the AC (signed) pixel values 808, 808'. The newly interpolated DC value 812 and the newly interpolated AC values 816 can then be combined 818, resulting in a new prediction 820 for the selected region.

[00128] As with the other interpolation values in this invention, the appropriate weightings can be signaled to a decoder per image unit. A small number of bits can select between a number of interpolation values, as well as selecting the independent interpolation of the AC versus DC aspects of the pixel values.

#### Linear & Non-Linear Interpolation

[00129] Interpolation is a linear weighted average. Since the interpolation operation is linear, and since the pixel values in each image frame are often represented in a non-linear form (such as video gamma or logarithmic representations), further optimization of the interpolation process becomes possible. For example, interpolation of pixels for a particular sequence of frames, as well as interpolation of DC values separately from AC values, will sometimes be optimal or near optimal with a linear pixel representation. However, for other frame sequences, such interpolation will be optimal or near optimal if the pixels are represented as logarithmic values or in other pixel representations. Further, the optimal or near optimal representations for interpolating U and V (chroma) signal components may differ from the optimal or near optimal representations for the Y (luminance) signal component. It is therefore a useful aspect of the invention to convert a pixel representation to an alternate representation as part of the interpolation procedure.

[00130] FIG. 9 is a flowchart showing one embodiment of a method for interpolation of luminance pixels using an alternative representation. Starting with a region or block of luminance (Y) pixels in an initial representation (e.g., video gamma or logarithmic) (Step 900), the pixel data is transformed to an alternative representation (e.g., linear, logarithmic, video gamma) different from the initial representation (Step 902). The transformed pixel region or block is then interpolated as described above (Step 906), and transformed back to the initial representation (Step 906). The result is interpolated pixel luminance values (Step 908).

[00131] FIG. 10 is a flowchart showing one embodiment of a method for interpolation of chroma pixels using an alternative representation. Starting with a region or block of chroma (U, V) pixels in an initial representation (e.g., video gamma or logarithmic) (Step

1000), the pixel data is transformed to an alternative representation (e.g., linear, logarithmic, video gamma) different from the initial representation (Step 1002). The transformed pixel region or block is then interpolated as described above (Step 1006), and transformed back to the initial representation (Step 1008). The result is interpolated pixel chroma values (Step 1008).

[00132] The transformations between representations may be performed in accordance with the teachings of U.S. Patent Application No. 09/905,039, entitled "Method and System for Improving Compressed Image Chroma Information", assigned to the assignee of the present invention and hereby incorporated by reference. Note that the alternative representation transformation and its inverse can often be performed using a simple lookup table.

[00133] As a variation of this aspect of the invention, the alternative (linear or non-linear) representation space for AC interpolation may differ from the alternative representation space for DC interpolation.

[00134] As with the interpolation weightings, the selection of which alternate interpolation representation is to be used for each of the luminance (Y) and chroma (U and V) pixel representations may be signaled to the decoder using a small number of bits for each selected image unit.

#### *Number of Motion Vectors per Macroblock*

[00135] In MPEG-2, one motion vector is allowed per 16x16 macroblock in P frames. In B frames, MPEG-2 allows a maximum of 2 motion vectors per 16x16 macroblock, corresponding to the bi-directional interpolative mode. In MPEG-4 video coding, up to 4 motion vectors are allowed per 16x16 macroblock in P frames, corresponding to one motion vector per 8x8 DCT block. In MPEG-4 B frames, a maximum of two motion vectors are allowed for each 16x16 macroblock, when using interpolative mode. A single motion vector delta in MPEG-4 direct mode can result in four independent "implicit" motion vectors, if the subsequent corresponding P frame macroblock was set in 8x8 mode having four vectors. This is achieved by adding the one motion vector delta carried in a 16x16 B frame macroblock to each of the corresponding four independent motion vectors from the following P frame macroblock, after scaling for the distance in time (the B frame is closer in time than the P frame's previous P or I frame reference).

[00136] One aspect of the invention includes the option to increase the number of motion vectors per picture region, such as a macroblock. For example, it will sometimes

prove beneficial to have more than two motion vectors per B frame macroblock. These can be applied by referencing additional P or I frames and having three or more interpolation terms in the weighted sum. Additional motion vectors can also be applied to allow independent vectors for the 8x8 DCT blocks of the B frame macroblock. Also, four independent deltas can be used to extend the direct mode concept by applying a separate delta to each of the four 8x8-region motion vectors from the subsequent P frame.

[00137] Further, P frames can be adapted using B-frame implementation techniques to reference more than one previous frame in an interpolative mode, using the B-frame two-interpolation-term technique described above. This technique can readily be extended to more than two previous P or I frames, with a resulting interpolation having three or more terms in the weighted sum.

[00138] As with other aspects of this invention (e.g., pixel representation and DC versus AC interpolation methods), particular weighted sums can be communicated to a decoder using a small number of bits per image unit.

[00139] In applying this aspect of the invention, the correspondence between 8x8 pixel DCT blocks and the motion vector field need not be as strict as with MPEG-2 and MPEG-4. For example, it may be useful to use alternative region sizes other than 16x16, 16x8 (used only with interlace in MPEG-4), and 8x8 for motion vectors. Such alternatives might include any number of useful region sizes, such as 4x8, 8x12, 8x16, 6x12, 2x8, 4x8, 24x8, 32x32, 24x24, 24x16, 8x24, 32x8, 32x4, etc. Using a small number of such useful sizes, a few bits can signal to a decoder the correspondence between motion vectors region sizes and DCT block sizes. In systems where a conventional 8x8 DCT block is used, a simple set of correspondences to the motion vector field are useful to simplify processing during motion compensation. In systems where the DCT block size is different from 8x8, then greater flexibility can be achieved in specifying the motion vector field, as described in co-pending U.S. Patent Application No. 09/545,233, entitled "Enhanced Temporal and Resolution Layering in Advanced Television", assigned to the assignee of the present invention and hereby incorporated by reference. Note that motion vector region boundaries need not correspond to DCT region boundaries. Indeed, it is often useful to define motion vector regions in such a way that a motion vector region edge falls within a DCT block (and not at its edge).

[00140] The concept of extending the flexibility of the motion vector field also applies to the interpolation aspect of this invention. As long as the correspondence between each pixel and one or more motion vectors to one or more reference frames is specified, the

interpolation method described above can be applied to the full flexibility of useful motion vectors using all of the generality of this invention. Even the size of the regions corresponding to each motion vector can differ for each previous frame reference when using P frames, and each previous and future frame reference when using B frames. If the region sizes for motion vectors differ when applying the improved interpolation method of this invention, then the interpolation reflects the common region of overlap. The common region of overlap for motion vector references can be utilized as the region over which the DC term is determined when separately interpolating DC and AC pixel values.

[00141] FIG. 11 is a diagram showing unique motion vector region sizes 1100, 1102 for each of two P frames 1104, 1106. Before computing interpolation values in accordance with this invention, the union 1108 of the motion vector region sizes is determined. The union 1108 defines all of the regions which are considered to have an assigned motion vector. [00142] Thus, for example, in interpolating 4x4 DCT regions of a B frame 1112 backwards to the prior P frame 1104, a 4x4 region 1110 within the union 1108 would use the motion vector corresponding to the 8x16 region 1114 in the prior P frame. If predicting forward, the region 1110 within the union 1108 would use the motion vector corresponding to the 4x16 region 1115 in the next P frame. Similarly, interpolation of the region 1116 within the union 1108 backwards would use the motion vector corresponding to the 8x16 region 1114, while predicting the same region forward would use the motion vector corresponding to the 12x16 region 1117.

[00143] In one embodiment of the invention, two steps are used to accomplish the interpolation of generalized (i.e., non-uniform size) motion vectors. The first step is to determine the motion vector common regions, as described with respect to FIG. 11. This establishes the correspondence between pixels and motion vectors (i.e., the number of motion vectors per specified pixel region size) for each previous or subsequent frame reference. The second step is to utilize the appropriate interpolation method and interpolation factors active for each region of pixels. It is a task of the encoder to ensure that optimal or near optimal motion vector regions and interpolation methods are specified, and that all pixels have their vectors and interpolation methods completely specified. This can be very simple in the case of a fixed pattern of motion vectors (such as one motion vector for each 32x8 block, specified for an entire frame), with a single specified interpolation method (such as a fixed proportion blend to each distance of referenced frame, specified for the entire frame). This method can also become quite complex if regional changes are made to the motion vector region sizes, and where the region sizes differ depending upon which previous or subsequent frame is

referenced (e.g., 8x8 blocks for the nearest previous frame, and 32x8 blocks for the next nearest previous frame). Further, the interpolation method may be regionally specified within the frame.

[00144] When encoding, it is the job of the encoder to determine the optimal or near optimal use of the bits to select between motion vector region shapes and sizes, and to select the optimal or near optimal interpolation method. A determination is also required to specify the number and distance of the frames referenced. These specifications can be determined by exhaustive testing of a number of candidate motion vector region sizes, candidate frames to reference, and interpolation methods for each such motion vector region, until an optimal or near optimal coding is found. Optimality (relative to a selected criteria) can be determined by finding the least SNR after encoding a block or the lowest number of bits for a fixed quantization parameter (QP) after coding the block, or by application of another suitable measure.

#### *Direct Mode Extension*

[00145] Conventional direct mode, used in B frame macroblocks in MPEG-4, can be efficient in motion vector coding, providing the benefits of 8x8 block mode with a simple common delta. Direct mode weights each corresponding motion vector from the subsequent P frame, which references the previous P frame, at the corresponding macroblock location based upon distance in time. For example, if M=3 (i.e., two intervening B frames), with simple linear interpolation the first B frame would use -2/3 times the subsequent P frame motion vector to determine a pixel offset with respect to such P frame, and 1/3 times the subsequent P frame motion vector to determine a pixel offset with respect to the previous P frame. Similarly, the second B frame would use -1/3 times the same P frame motion vector to determine a pixel offset with respect to such P frame, and 2/3 times the subsequent P frame motion vector to determine a pixel offset with respect to the previous P frame. In direct mode, a small delta is added to each corresponding motion vector. As another aspect of this invention, this concept can be extended to B frame references which point to one or more n-away P frames, which in turn reference one or more previous or subsequent P frames or l frames, by taking the frame distance into account to determine a frame scale fraction. [00146] FIG. 12 is a diagram showing a sequence of P and B frames with interpolation weights for the B frames determined as a function of distance from a 2-away subsequent P frame that references a 1-away subsequent P frame. In the illustrated example, M=3, indicating two consecutive B frames 1200, 1202 between bracketing P frames 1204, 1206. In



this example, each co-located macroblock in the next nearest subsequent P frame 1208 (i.e.,  $n=2$ ) might point to the intervening (i.e., nearest) P frame 1204, and the first two B frames 1200, 1202 may reference the next nearest subsequent P frame 1208 rather than the nearest subsequent P frame 1204, as in conventional MPEG. Thus, for the first B frame 1200, the frame scale fraction  $5/3$  times the motion vector  $mv$  from the next nearest subsequent P frame 1208 would be used as a pixel offset with respect to P frame 1208, and the second B frame 1202 would use an offset of  $4/3$  times that same motion vector.

[00147] If a nearest subsequent P frame referenced by a B frame points to the next nearest previous P frame, then again the simple frame distance can be used to obtain the suitable frame scale fraction to apply to the motion vectors. FIG. 13 is a diagram showing a sequence of P and B frames with interpolation weights for the B frames determined as a function of distance from a 1-away subsequent P frame that references a 2-away previous P frame. In the illustrated example,  $M=3$ , and B frames 1300, 1302 reference the nearest subsequent P frame 1304, which in turn references the 2-away P frame 1306. Thus, for the first B frame 1300, the pixel offset fraction is the frame scale fraction  $2/6$  multiplied by the motion vector  $mv$  from the nearest subsequent P frame 1304, and the second B frame 1302 would have a pixel offset of the frame scale fraction  $1/6$  multiplied by that same motion vector, since the motion vector of the nearest subsequent P frame 1304 points to the 2-away previous P frame 1306, which is 6 frames distant.

[00148] In general, in the case of a B frame referencing a single P frame in direct mode, the frame distance method sets the numerator of a frame scale fraction equal to the frame distance from that B frame to its referenced, or "target", P frame, and sets the denominator equal to the frame distance from the target P frame to another P frame referenced by the target P frame. The sign of the frame scale fraction is negative for measurements made from a B frame to a subsequent P frame, and positive for measurements made from a B frame to a prior P frame. This simple method of applying a frame-distance or the frame scale fraction to a P frame motion vector can achieve an effective direct mode coding.

[00149] Further, another aspect of this invention is to allow direct mode to apply to multiple interpolated motion vector references of a P frame. For example, if a P frame was interpolated from the nearest and next nearest previous P frames, direct mode reference in accordance with this aspect of the invention allows an interpolated blend for each multiple reference direct mode B frame macroblock. In general, the two or more motion vectors of a P frame can have an appropriate frame scale fraction applied. The two or more frame-distance

modified motion vectors then can be used with corresponding interpolation weights for each B frame referencing or targeting that P frame, as described below, to generate interpolated B frame macroblock motion compensation.

[00150] FIG. 14 is a diagram showing a sequence of P and B frames in which a subsequent P frame has multiple motion vectors referencing prior P frames. In this example, a B frame 1400 references a subsequent P frame P3. This P3 frame in turn has two motion vectors,  $mv1$  and  $mv2$ , that reference corresponding prior P frames P2, P1. In this example, each macroblock of the B frame 1400 can be interpolated in direct mode using either of two weighting terms or a combination of such weighting terms.

[00151] Each macroblock for the B frame 1400 would be constructed as a blend from:

- [00152] • corresponding pixels of frame P2 displaced by the frame scale fraction  $1/3$  of  $mv1$  (where the pixels may then be multiplied by some proportional weight  $l$ ) plus corresponding pixels of frame P3 displaced by the frame scale fraction  $-2/3$  of  $mv1$  (where the pixels may then be multiplied by some proportional weight  $l$ ); and
- [00153] • corresponding pixels of frame P1 displaced by the frame scale fraction  $2/3$  ( $4/6$ ) of  $mv2$  (where the pixels may then be multiplied by some proportional weight  $k$ ) plus corresponding pixels of frame P3 displaced by the frame scale fraction  $-1/3$  ( $-2/6$ ) of  $mv2$  (where the pixels may then be multiplied by some proportional weight  $l$ ).

[00154] As with all direct modes, a motion vector delta can be utilized with each of  $mv1$  and  $mv2$ .

[00155] In accordance with this aspect of the invention, direct mode predicted macroblocks in B frames can also reference multiple subsequent P frames, using the same methodology of interpolation and motion vector frame scale fraction application as with multiple previous P frames. FIG. 15 is a diagram showing a sequence of P and B frames in which a nearest subsequent P frame has a motion vector referencing a prior P frame, and a next nearest subsequent P frame has multiple motion vectors referencing prior P frames. In this example, a B frame 1500 references two subsequent P frames P2, P3. The P3 frame has two motion vectors,  $mv1$  and  $mv2$ , that reference corresponding prior P frames P2, P1. The P2 frame has one motion vector,  $mv3$ , which references the prior P frame P1. In this example, each macroblock of the B frame 1500 is interpolated in direct mode using three weighting



terms. In this case, the motion vector frame scale fractions may be greater than 1 or less than -1.

[00156] The weightings for this form of direct mode B frame macroblock interpolation can utilize the full generality of interpolation as described herein. In particular, each weight, or combinations of the weights, can be tested for best performance (e.g., quality versus number of bits) for various image units. The interpolation fraction set for this improved direct mode can be specified to a decoder with a small number of bits per image unit.

[00157] Each macroblock for the B frame 1500 would be constructed as a blend from:

- [00158] • corresponding pixels of frame P3 displaced by the frame scale fraction  $-5/3$  of  $mv1$  (where the pixels may then be multiplied by some proportional weight  $l$ ) plus corresponding pixels of frame P2 displaced by the frame scale fraction  $-2/3$  of  $mv1$  (where the pixels may then be multiplied by some proportional weight  $j$ );

- [00159] • corresponding pixels of frame P3 displaced by the frame scale fraction  $-5/6$  of  $mv2$  (where the pixels may then be multiplied by some proportional weight  $k$ ) plus corresponding pixels of frame P1 displaced by the frame scale fraction  $1/6$  of  $mv2$  (where the pixels may then be multiplied by some proportional weight  $i$ ); and

- [00160] • corresponding pixels of frame P2 displaced by the frame scale fraction  $-2/3$  of  $mv3$  (where the pixels may then be multiplied by some proportional weight  $m$ ) plus corresponding pixels of frame P1 displaced by the frame scale fraction  $1/3$  of  $mv3$  (where the pixels may then be multiplied by some proportional weight  $n$ ).

[00161] As with all direct modes, a motion vector delta can be utilized with each of  $mv1$ ,  $mv2$ , and  $mv3$ .

[00162] Note that a particularly beneficial direct coding mode often occurs when the next nearest subsequent P frame references the nearest P frames bracketing a candidate B frame.

[00163] Direct mode coding of B frames in MPEG-4 always uses the subsequent P frame's motion vectors as a reference. In accordance with another aspect of the invention, it is also possible for a B frame to reference the motion vectors of the *previous* P frame's co-located macroblocks, which will sometimes prove a beneficial choice of direct mode coding reference. In this case, the motion vector frame scale fractions will be greater than one, when

the next nearest previous P frame is referenced by the nearest previous frame's P motion vector. FIG. 16 is a diagram showing a sequence of P and B frames in which a nearest previous P frame has a motion vector referencing a prior P frame. In this example, a B frame 1600 references the 1-away previous P frame P2. The motion vector  $mv$  of frame P2 references the next previous P frame P1 (2-away relative to the B frame 1600). The appropriate frame scale fractions are shown.

[00164] If the nearest previous P frame is interpolated from multiple vectors and frames, then methods similar to those described in conjunction with FIG. 14 apply to obtain the motion vector frame scale fractions and interpolation weights. FIG. 17 is a diagram showing a sequence of P and B frames in which a nearest previous P frame has two motion vectors referencing prior P frames. In this example, a B frame 1700 references the previous P frame P3. One motion vector  $mv1$  of the previous P3 frame references the next previous P frame P2, while the second motion vector  $mv2$  references the 2-away previous P frame P1. The appropriate frame scale fractions are shown.

[00165] Each macroblock for the B frame 1700 would be constructed as a blend from:

- [00166] • corresponding pixels of frame P3 displaced by the frame scale fraction  $1/3$  of  $mv1$  (where the pixels may then be multiplied by some proportional weight  $l$ ) plus corresponding pixels of frame P2 displaced by the frame scale fraction  $4/3$  of  $mv1$  (where the pixels may then be multiplied by some proportional weight  $j$ ); and
- [00167] • corresponding pixels of frame P3 displaced by the frame scale fraction  $1/6$  of  $mv2$  (where the pixels may then be multiplied by some proportional weight  $k$ ) plus corresponding pixels of frame P1 displaced by the frame scale fraction  $7/6$  of  $mv2$  (where the pixels may then be multiplied by some proportional weight  $i$ ).

[00168] When the motion vector of a previous P frame (relative to a B frame) points to the next nearest previous P frame, it is not necessary to only utilize the next nearest previous frame as the interpolation reference, as in FIG. 16. The nearest previous P frame may prove a better choice for motion compensation. In this case, the motion vector of the nearest previous P frame is shortened to the frame distance fraction from a B frame to that P frame. FIG. 18 is a diagram showing a sequence of P and B frames in which a nearest previous P frame has a motion vector referencing a prior P frame. In this example, for  $M=3$ , a first B frame 1800 would use  $1/3$  and  $-2/3$  frame distance fractions times the motion vector  $mv$  of the nearest

previous P frame P2. The second B frame 1802 would use  $2/3$  and  $-1/3$  frame distance fractions (not shown). Such a selection would be signaled to the decoder to distinguish this case from the case shown in FIG. 16.

**[00169]** As with all other coding modes, the use of direct mode preferably involves testing the candidate mode against other available interpolation and single-vector coding modes and reference frames. For direct mode testing, the nearest subsequent P frame (and, optionally, the next nearest subsequent P frame or even more distant subsequent P frames, and/or one or more previous P frames) can be tested as candidates, and a small number of bits (typically one or two) can be used to specify the direct mode P reference frame distance(s) to be used by a decoder.

#### *Extended Interpolation Values*

**[00170]** It is specified in MPEG-1, 2, and 4, as well as in the H.261 and H.263 standards, that B frames use an equal weighting of pixel values of the forward referenced and backward referenced frames, as displaced by the motion vectors. Another aspect of this invention includes application of various useful unequal weightings that can significantly improve B frame coding efficiency, as well as the extension of such unequal weightings to more than two references, including two or more references backward or forward in time. This aspect of the invention also includes methods for more than one frame being referenced and interpolated for P frames. Further, when two or more references point forward in time, or when two or more references point backward in time, it will sometimes be useful to use negative weightings as well as weightings in excess of 1.0.

**[00171]** For example, FIG. 19 is a frame sequence showing the case of three P frames P1, P2, and P3, where P3 uses an interpolated reference with two motion vectors, one for each of P1 and P2. If, for example, a continuous change is occurring over the span of frames between P1 and P3, then P2-P1 (i.e., the pixel values of frame P2, displaced by the motion vector for P2, minus the pixel values of frame P1, displaced by the motion vector for P1) will equal P3-P2. Similarly, P3-P1 will be double the magnitude of P2-P1 and P3-P2. In such a case, the pixel values for frame P3 can be predicted differentially from P1 and P2 through the formula:

$$P3 = P1 + 2 \times (P2 - P1) = (2 \times P2) - P1$$

**[00172]** In this case, the interpolative weights for P3 are 2.0 for P2, and -1.0 for P1.

**[00173]** As another example, FIG. 20 is a frame sequence showing the case of four P frames P1, P2, P3, and P4, where P4 uses an interpolated reference with three motion vectors,

one for each of P1, P2, and P3. Thus, since P4 is predicted from P3, P2, and P1, three motion vectors and interpolative weights would apply. If, in this case, a continuous change were occurring over this span of frames, then P2-P1 would equal both P3-P2 and P4-P3, and P4-P1 would equal both  $3 \times (P2-P1)$  and  $3 \times (P3-P2)$ .

**[00174]** Thus, in this example case, a prediction of P4 based upon P2 and P1 would be:  
 $P4 = P1 + 3 \times (P2 - P1) = (3 \times P2) - (2 \times P1)$  (weights 3.0 and -2.0)

**[00175]** The prediction of P4 based upon P3 and P1 would be:

$$P4 = P1 + 3/2 \times (P3 - P1) = (3/2 \times P3) - (1/2 \times P1) \quad (\text{weights } 1.5 \text{ and } -0.5)$$

**[00176]** The prediction of P4 based upon P3 and P2 would be:

$$P4 = P2 + 2 \times (P3 - P2) = (2 \times P3) - P2 \quad (\text{weights } 2.0 \text{ and } -1.0)$$

**[00177]** However, it might also be likely that the change most near to P4, involving P3 and P2, is a more reliable predictor of P4 than predictions involving P1. Thus, by giving  $1/4$  weight to each of the two terms above involving P1, and  $1/2$  weight to the term involving only P3 and P2, would result in:

$$\begin{aligned} &1/2(2 P3 - P2) + 1/4(3/2 P3 - 1/2 P1) + 1/4(3 P2 - 2 P1) = \\ &1 \ 3/8 P3 + 1/4 P2 - 5/8 P1 \quad (\text{weights } 1.375, 0.25, \text{ and } -0.625) \end{aligned}$$

**[00178]** Accordingly, it will sometimes be useful to use weights both above 1.0 and below zero. At other times, if there is noise-like variation from one frame to the next, a positive weighted average having mild coefficients between 0.0 and 1.0 might yield the best predictor of P4's macroblock (or other region of pixels). For example, an equal weighting of  $1/3$  of each of P1, P2, and P3 in FIG. 20 might form the best predictor of P4 in some cases.

**[00179]** Note that the motion vector of the best match is applied to determine the region of P1, P2, P3, etc., being utilized by the computations in this example. This match might best be an AC match in some cases, allowing a varying DC term to be predicted through the AC coefficients. Alternatively, if a DC match (such as Sum of Absolute Difference) is used, then changes in AC coefficients can often be predicted. In other cases, various forms of motion vector match will form a best prediction with various weighting blends. In general, the best predictor for a particular case is empirically determined using the methods described herein.

**[00180]** These techniques are also applicable to B frames that have two or more motion vectors pointing either backward or forward in time. When pointing forward in time, the coefficient pattern described above for P frames is reversed to accurately predict backward to the current P frame. It is possible to have two or more motion vectors in both the forward and backward direction using this aspect of the invention, thereby predicting in both

directions concurrently. A suitable weighted blend of these various predictions can be optimized by selecting the blend weighting which best predicts the macroblock (or other pixel region) of a current B frame.

[00181] FIG. 21 is a diagram showing a sequence of P and B frames in which various P frames have one or more motion vectors referencing various previous P frames, and showing different weights  $\alpha$ -e assigned to respective forward and backward references by a particular B frame. In this example, a B frame 2100 references three previous P frames and two subsequent P frames.

[00182] In the example illustrated in FIG. 21, frame P5 must be decoded for this example to work. It is useful sometimes to order frames in a bitstream in the order needed for decoding ("delivery order"), which is not necessarily the order of display ("display order"). For example, in a frame sequence showing cyclic motion (e.g., rotation of an object), a particular P frame may be more similar to a distant P frame than to the nearest subsequent P frame. FIG. 22 is a diagram showing a sequence of P and B frames in which the bitstream delivery order of the P frames differs from the display order. In this example, frame P3 is more similar to frame P5 than to frame P4. It is therefore useful to deliver and decode P5 before P4, but display P4 before P5. Preferably, each P frame should signal to the decoder when such P frame can be discarded (e.g., an expiration of  $n$  frames in bitstream order, or after frame X in the display order).

[00183] If the weightings are selected from a small set of choices, then a small number of bits can signal to the decoder which weighting is to be used. As with all other weightings described herein, this can be signaled to a decoder once per image unit, or at any other point in the decoding process where a change in weightings is useful.

[00184] It is also possible to download new weightings sets. In this way, a small number of weighting sets may be active at a given time. This allows a small number of bits to signal a decoder which of the active weighting sets is to be used at any given point in the decoding process. To determine suitable weighting sets, a large number of weightings can be tested during encoding. If a small subset is found to provide high efficiency, then that subset can be signaled to a decoder for use. A particular element of the subset can thus be signaled to the decoder with just a few bits. For example, 10 bits can select 1 of 1024 subset elements. Further, when a particular small subset should be changed to maintain efficiency, a new subset can be signaled to the decoder. Thus, an encoder can dynamically optimize the number of bits required to select among weighting set elements versus the number of bits needed to update the weighting sets. Further, a small number of short codes can be used to signal

common useful weightings, such as  $1/2$ ,  $1/3$ ,  $1/4$ , etc. In this way, a small number of bits can be used to signal the set of weightings, such as for a K-forward-vector prediction in a P frame (where  $K = 1, 2, 3, \dots$ ), or a K-forward-vector and L-backward-vector prediction in a B frame (where K and L are selected from  $0, 1, 2, 3, \dots$ ), or a K-forward-vector and L-backward-vector prediction in a P frame (where K and L are selected from  $0, 1, 2, 3, \dots$ ), as a function of the current M value (i.e., the relative position of the B frame with respect to the neighboring P (or I) frames).

[00185] FIG. 23 is a diagram showing a sequence of P and B frames with assigned weightings. A B frame 2300 has weights  $\alpha$ -e, the values of which are assigned from a table of B frame weighting sets 2302. A P frame 2304 has weights  $m$ ,  $n$ , the values of which are assigned from a table of P frame weighting sets 2306. Some weightings can be static (i.e., permanently downloaded to the decoder), and signaled by an encoder. Other weightings may be dynamically downloaded and then signaled.

[00186] This same technique may be used to dynamically update weighting sets to select DC interpolation versus AC interpolation. Further, code values can be signaled which select normal (linear) interpolation (of pixel values normally represented in a non-linear representation) versus linear interpolation of converted values (in an alternate linear or non-linear representation). Similarly, such code values can signal which such interpolation to apply to AC or DC values or whether to split AC and DC portions of the prediction.

[00187] Active subsetting can also be used to minimize the number of bits necessary to select between the sets of weighting coefficients currently in use. For example, if 1024 downloaded weighting sets were held in a decoder, perhaps 16 might need to be active during one particular portion of a frame. Thus, by selecting which subset of 16 (out of 1024) weighting sets are to be active, only 4 bits need be used to select which weighting set of these 16 is active. The subsets can also be signaled using short codes for the most common subsets, thus allowing a small number of bits to select among commonly used subsets.

#### *Softening and Sharpening*

[00188] As with the simple separation of a DC component from AC signals via subtraction of the average value, other filtering operations are also possible during motion vector compensated prediction. For example, various high-pass, band-pass, and low-pass filters can be applied to a pixel region (such as a macroblock) to extract various frequency bands. These frequency bands can then be modified when performing motion compensation. For example, it often might be useful on a noisy moving image to filter out the highest

frequencies in order to soften (make less sharp, or blur slightly) the image. The softer image pixels, combined with a steeper tilt matrix for quantization (a steeper tilt matrix ignores more high-frequency noise in the current block), will usually form a more efficient coding method. It is already possible to signal a change in the quantization tilt matrix for every image unit. It is also possible to download custom tilt matrices for luminance and chroma. Note that the effectiveness of motion compensation can be improved whether the tilt matrix is changed or not. However, it will often be most effective to change both the tilt matrix and filter parameters which are applied during motion compensation.

**[00189]** It is common practice to use reduced resolution for chroma coding together with a chroma specific tilt matrix. However, the resolution of chroma coding is static in this case (such as 4:2:0 coding half resolution vertically and horizontally, or 4:2:2 coding half resolution only horizontally). Coding effectiveness can be increased in accordance with this aspect of the invention by applying a dynamic filter process during motion compensation to both chroma and luminance (independently or in tandem), selected per image unit.

**[00190]** U.S. Patent Application No. 09/545,233, entitled "Enhanced Temporal and Resolution Layering in Advanced Television" (referenced above), describes the use of improved displacement filters having negative lobes (a truncated sinc function). These filters have the advantage that they preserve sharpness when performing the fractional-pixel portion of motion vector displacement. At both the integer pixel displacement point and at the fractional points, some macroblocks (or other useful image regions) are more optimally displaced using filters which reduce or increase their sharpness. For example, for a "back focus" (where some objects in the frame are going out of focus over time, and others portions of the frame are coming into focus), the transition is one of change both in sharpness and in softness. Thus, a motion compensation filter that can both increase sharpness at certain regions in an image while decreasing sharpness in other regions can improve coding efficiency. In particular, if a region of a picture is going out of focus, it may be beneficial to decrease sharpness, which will soften the image (thereby potentially creating a better match) and decrease grain and/or noise (thereby possibly improving coding efficiency). If a region of the image is coming into focus, it may be beneficial to preserve maximum sharpness, or even increase sharpness using larger negative lobe filter values.

**[00191]** Chroma filtering can also benefit from sharpness increase and decrease during coding. For example, much of the coding efficiency benefits of 4:2:0 coding (half resolution chroma horizontally and vertically) can be achieved by using softer motion compensation filters for chroma while preserving full resolution in the U and/or V channels. Only when

color detail in the U and V channels is high will it be necessary to select the sharpest displacement filters; softer filters will be more beneficial where there is high color noise or grain.

**[00192]** In addition to changes in focus, it is also common to have the direction and amount of motion blur change from one frame to the next. At the motion picture film frame rate of 24 fps, even a simple dialog scene can have significant changes in motion blur from one frame to the next. For example, an upper lip might blur in one frame, and sharpen in the next, entirely due to the motion of the lip during the open shutter time in the camera. For such motion blur, it will be beneficial not only to have sharpening and softening (blurring) filters during motion compensation, but also to have a directional aspect to the sharpening and softening. For example, if a direction of motion can be determined, a softening or sharpening along that direction can be used to correspond to the moving or stopping of an image feature. The motion vectors used for motion compensation can themselves provide some useful information about the amount of motion, and the change in the amount of motion (*i.e.*, motion blur), for a particular frame (or region within a frame) with respect to any of the surrounding frames (or corresponding regions). In particular, a motion vector is the best movement match between P frames, while motion blur results from movement during the open shutter time within a frame.

**[00193]** FIG. 24 is a graph of position of an object within a frame versus time. The shutter of a camera is open only during part of a frame time. Any motion of the object while the shutter is open results in blur. The amount of motion blur is indicated by the amount of position change during the shutter open time. Thus, the slope of the position curve 2400 while the shutter is open is a measurement of motion blur.

**[00194]** The amount of motion blur and the direction of motion can also be determined from a combination of sharpness metrics, surrounding motion vectors (where image regions match), feature smear detection, and human assisted designation of frame regions. A filter can be selected based on the determined amount of motion blur and motion direction. For example, a mapping of various filters versus determined motion blur and direction can be empirically determined.

**[00195]** When combined with the other aspects of this invention, such intelligently applied filters can significantly improve compression coding efficiency. A small number of such filters can be selected with a small number of bits signaled to the decoder. Again, this can be done once per image unit or at other useful points in the decoding process. As with weighting sets, a dynamically loaded set of filters can be used, as well as an active subsetting

mechanism, to minimize the number of bits needed to select between the most beneficial set of filter parameters.

#### *Implementation*

[00196] The invention may be implemented in hardware or software, or a combination of both (e.g., programmable logic arrays). Unless otherwise specified, the algorithms included as part of the invention are not inherently related to any particular computer or other apparatus. In particular, various general purpose machines may be used with programs written in accordance with the teachings herein, or it may be more convenient to construct more specialized apparatus (e.g., integrated circuits) to perform particular functions. Thus, the invention may be implemented in one or more computer programs executing on one or more programmable computer systems each comprising at least one processor, at least one data storage system (including volatile and non-volatile memory and/or storage elements), at least one input device or port, and at least one output device or port. Program code is applied to input data to perform the functions described herein and generate output information. The output information is applied to one or more output devices, in known fashion.

[00197] Each such program may be implemented in any desired computer language (including machine, assembly, or high level procedural, logical, or object oriented programming languages) to communicate with a computer system. In any case, the language may be a compiled or interpreted language.

[00198] Each such computer program is preferably stored on or downloaded to a storage media or device (e.g., solid state memory or media, or magnetic or optical media) readable by a general or special purpose programmable computer, for configuring and operating the computer when the storage media or device is read by the computer system to perform the procedures described herein. The inventive system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer system to operate in a specific and predefined manner to perform the functions described herein.

[00199] A number of embodiments of the invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. For example, some of the steps described above may be order independent, and thus can be performed in an order different from that described. Accordingly, other embodiments are within the scope of the following claims.

#### **WHAT IS CLAIMED IS:**

1. A method for video image compression, the method comprising: providing a sequence of frames including picture regions, the frames including predicted frames and referenceable frames; and encoding a picture region of at least one predicted frame by reference to two or more prior referenceable frames in the sequence.
2. The method of claim 1, wherein said at least one predicted frame comprises a referenceable frame.
3. The method of claim 1, wherein said referenceable frames comprise predicted frames and intra frames.
4. The method of claim 1, wherein said predicted frames comprise referenceable predicted frames and bidirectional predicted frames.
5. The method of claim 4, wherein said at least one predicted frame comprises a bidirectional predicted frame.
6. The method of claim 5, further comprising: wherein encoding a picture region of a bidirectional predicted frame by reference to two or more subsequent referenceable frames in the sequence.
7. The method of claim 1, further comprising: identifying at least one of said two or more previous referenceable frames.
8. The method of claim 1, wherein said encoding comprises encoding using an unequal weighting of selected picture regions from said two or more previous referenceable frames.
9. The method of claim 8, further comprising: identifying said unequal weighting.

10. The method of claim 8, wherein the unequal weighting includes weights greater than one or less than zero.

11. The method of claim 8, wherein the frames comprise frames arranged in picture regions, and wherein said encoding comprises encoding using unequal pixel values corresponding to said two or more referenceable frames.

12. The method of claim 1, wherein the sequence of frames include referenceable frames and bidirectional predicted frames, each of said frames including pixel values arranged in macroblocks.

13. The method of claim 12, further comprising: determining at least one macroblock within a bidirectional predicted frame using direct mode prediction based on motion vectors from one or more predicted frames in display order.

14. The method of claim 13, wherein said encoding comprises encoding using an unequal weighting of selected picture regions from said two or more previous referenceable frames.

15. The method of claim 14, wherein at least one such motion vector is scaled by a frame scale fraction of less than zero or greater than one.


16. A method for video image compression, the method comprising: providing a sequence of frames including picture regions, the frames including bidirectional predicted (B) frames and referenceable frames including predicted (P) frames and intra (I) frames; and encoding a picture region of at least one P frame by reference to a prior referenceable frame in the sequence, wherein said prior referenceable frame is spaced from said P frame by at least one intervening referenceable frame.

17. The method of claim 16, further comprising: identifying the prior referenceable frame; and signaling a decoder with said identification.


18. A method, in a video image compression system having a sequence of referenceable and bidirectional predicted frames, for dynamically determining a code pattern of such frames having a variable number of bidirectional predicted frames, the method including:

selecting an initial sequence beginning with a referenceable frame, having at least one immediately subsequent bidirectional predicted frame, and ending in a referenceable frame; adding a referenceable frame to the end of the initial sequence to create a test sequence; evaluating the test sequence against a selected evaluation criteria; for each satisfactory step of evaluating the test sequence, inserting a bidirectional frame before the added referenceable frame and repeating the step of evaluating; and if the step of evaluating the test sequence is unsatisfactory, then accepting the prior test sequence as a current code pattern.


19. The method of claim 18, further comprising: determining an interpolation blend proportion for at least one picture region of the frames of a selected test sequence.



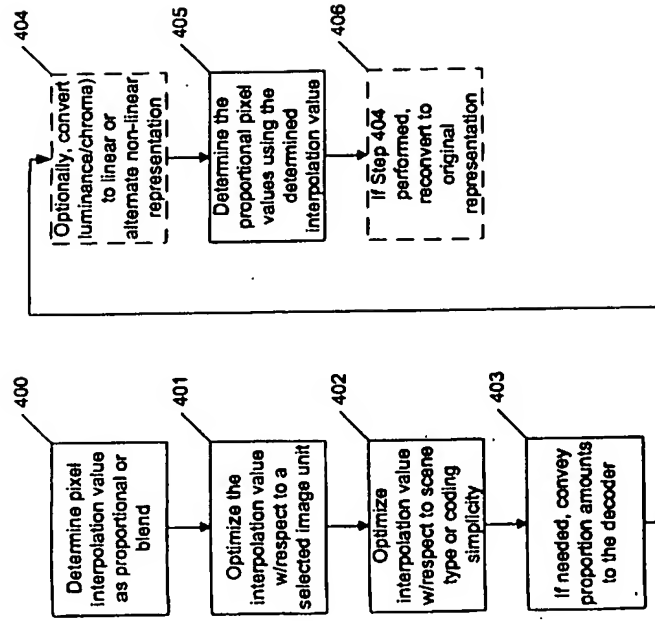
**FIG. 1**  
Proportional Motion Vector Weighting  
(Prior Art)




**FIG. 2**  
Pixel Value Proportional Weighting




**FIG. 3**  
Blended Pixel Value Weighting




**FIG. 4**



**FIG. 5**



**FIG. 6A**  
**Prior Art**



**FIG. 6B**  
**Prior Art**



188 P 700  
 188 P 702  
 188 P 704  
 188 P 706  
 188 P 708  
 188 P 710  
 188 P 712

FIG. 7

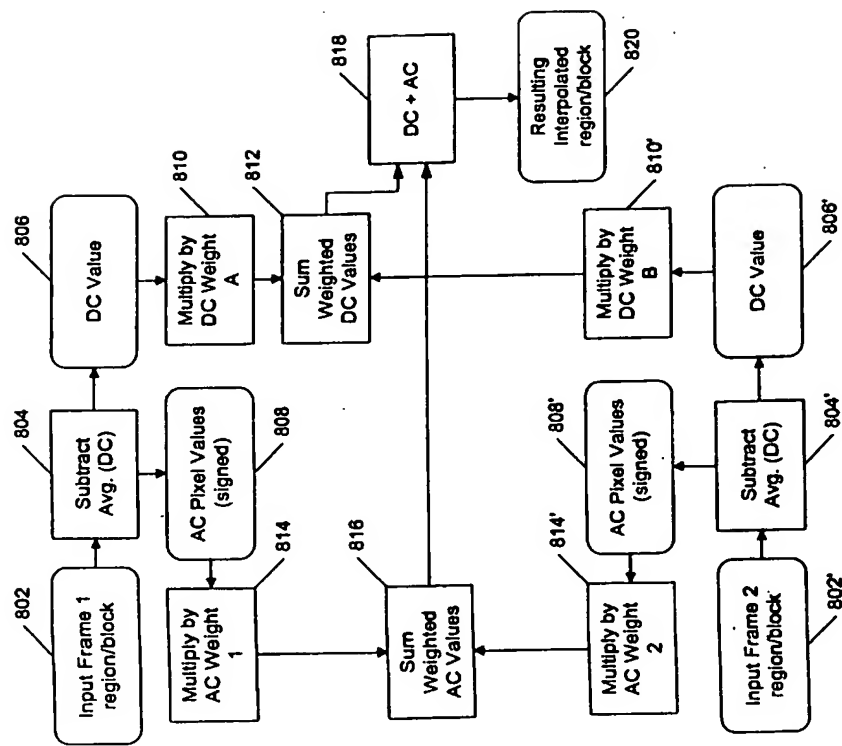





FIG. 8



**FIG. 11**






FIG. 12




FIG. 13




FIG. 14




FIG. 15




FIG. 16




FIG. 18




FIG. 17




FIG. 19




FIG. 20




FIG. 21




FIG. 22

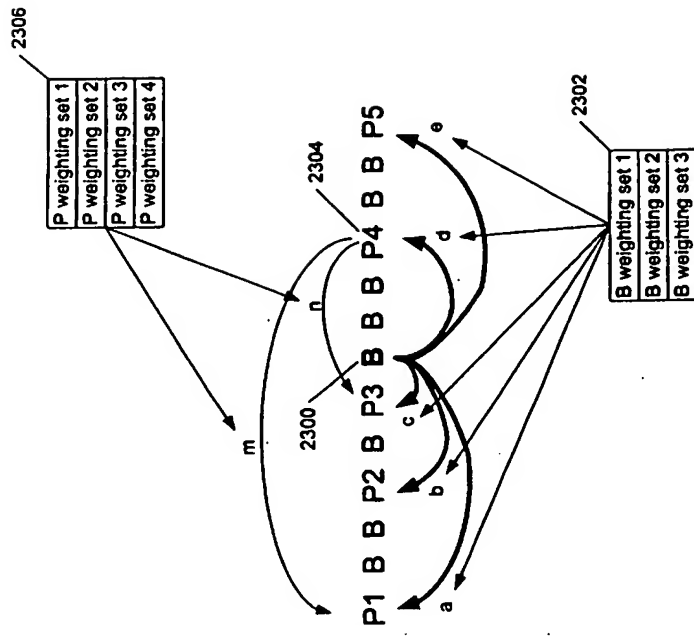


FIG. 23




FIG. 24